NASnews

September - October 1998 • Volume 3, Number 5

In this issue...

Feiereisen Sees NAS Future as Distributed Computing Hub

by Wade Roush

On July 20, after working under acting leadership for several years, staff in the NAS Systems Division welcomed their new permanent division chief, William J. Feiereisen.

It was a day Feiereisen says he had long anticipated. At heart, the 46-year-old Wisconsin native is a computational fluid dynamics (CFD) researcher with three degrees in mechanical engineering and a special interest in turbulence simulations and Monte Carlo simulations of high-speed aerodynamics problems. But in 1991, Feiereisen put aside his scientific work for a seven-year tour of duty in NASA program management, most recently in the High Performance Computing and Communications (HPCC) Program.

"I agreed to serve as a program manager," he says. "But the NAS Division job gives me the opportunity to get much closer to the technical matters, which is where I started out from and where I want to return."

To The Article...

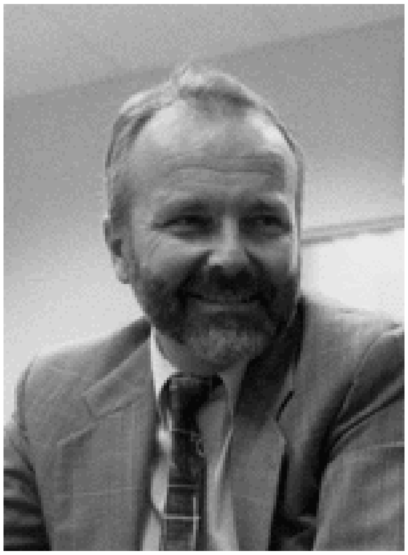

William J. Feiereisen. *Photo by Judy Conlon.*

▶ Main Menu

Feiereisen Sees NAS Future as Distributed Computing Hub

by Wade Roush

On July 20, after working under acting leadership for several years, staff in the NAS Systems Division welcomed their new permanent division chief, William J. Feiereisen.

It was a day Feiereisen says he had long anticipated. At heart, the 46-year-old Wisconsin native is a computational fluid dynamics (CFD) researcher with three degrees in mechanical engineering and a special interest in turbulence simulations and Monte Carlo simulations of high-speed aerodynamics problems. But in 1991, Feiereisen put aside his scientific work for a seven-year tour of duty in NASA program management, most recently in the High Performance Computing and Communications (HPCC) Program.

"I agreed to serve as a program manager," he says. "But the NAS Division job gives me the opportunity to get much closer to the technical matters, which is where I started out from and where I want to return."

Feiereisen's enthusiasm is also reflected in his plans for the future of the division. He says he wants to use the division chief position to help shape NAS as a major intellectual and technical hub for the young discipline of distributed heterogeneous computing.

Researchers working in this field want to make more efficient use of the nation's high-performance computing resources by connecting academic and government supercomputing centers into a seamless, nationwide "Information Power Grid" or IPG (see story, this issue). Feiereisen believes that as NAS takes the lead in building the IPG, the exciting ambience at Ames in the late '70s and early '80s, when Feiereisen was a graduate student at nearby Stanford University, can be recaptured.

'A Stimulating Atmosphere'

As the IPG effort begins, Feiereisen will put to use his broad experience in both engineering and management. Engineering seems to run in the family: Feiereisen's father was an aerospace engineer who in the 1960s settled at the University of Wisconsin as a professor of mechanical engineering. He set such an example that Feiereisen and his three brothers went on to earn a total of 10 degrees in the same field.

Feiereisen finished his first degree at Wisconsin in 1975, then moved to Stanford, where he did doctoral work until 1981 under turbulence pioneer Bill Reynolds. At that time, Ames was "one of the most exciting places on Earth to work," says Feiereisen. "They had the best [CFD] experts in the world, the best supercomputing people, they had an agency that was all behind them, and people were just coming here because it was the center of the world." Among those whom Feiereisen worked with were Harvard Lomax, Richard Waring, and Robert MacCormack, all well known in the CFD modeling world.

Feiereisen completed his doctoral work on numerical simulation of turbulent flows in 1981 and decided to see how he could apply this work in industry. He moved to Baden, Switzerland, spending several years with Brown-Boveri Corp. , a manufacturer of large steam, gas, and water turbines (now part of the giant industrial conglomerate ABB). Soon, however, Feiereisen began to miss what he calls "a more highly stimulating intellectual atmosphere" in the United States. In 1986, he returned to Moffett Field to take a position in the NASA Ames Aerothermodynamics Branch, where he developed simulations of the aerodynamics and air chemistry around aerospace vehicles reentering the atmosphere.

In 1991, management called: James Arnold, chief of the Ames Space Technology Division, suggested a temporary assignment to NASA headquarters so that Feiereisen could see how the agency operates in Washington. "That would arm me to deal with some of the things I needed to handle when I came back," he explained. Soon after his return to Ames, he became acting chief of the Aerothermodynamics Branch.

In 1993, Ames Aeronautics Office director F. Ron Bailey, the founding chief of the NAS Systems Division, asked Feiereisen to become project manager for the Computational Aerosciences wing of HPCC. Then in 1996, when NASA headquarters spun off management of many technical programs to the various NASA centers, HPCC migrated to Ames, and Feiereisen became overall program manager.

Partnership is Key

While he was with HPCC, Feiereisen recounts, two key trends emerged in the field of supercomputing. First, rising costs were forcing many aerospace companies to reduce their reliance on big, powerful, centralized machines such as those made by Cray Research. Second, many companies, academic institutions, and government agencies began to see that distributed computing-tying together the operations of multiple computers at local and distant sites-could both lower costs and increase efficiency. "You could have something that's larger and more powerful and more capable than anything that you could have at one specific site," Feiereisen explains.

Now is the time for Ames to act on this vision, Feiereisen says. Two university-led consortia, one based at the National Center for Supercomputing Applications at the University of Illinois and the other at the University of California at San Diego, have already won major support from the National Science Foundation for research on distributed supercomputing techniques. One goal of the consortia, for example, is to make the grid "transparent," by developing administrative and scheduling software that distributes computational jobs around the network so smoothly that users don't have to concern themselves with where and how their jobs will run. The NSF wanted another partner with systems engineering experience to implement such software and test it using real-world applications, Feiereisen says. "NAS will apply its considerable engineering and integration expertise to make this happen in concert with the NSF," he says.

Sharing the Vision

If the IPG project unfolds as Feiereisen envisions it, NAS will continue to be a center for supercomputing expertise in the aerospace industry and will also become a systems integration site for the prototype grid, with much research focused on hardware and software issues related to distributed computing. He has already begun to share this vision in informal get-acquainted sessions with NAS groups. In a NAS Technical Forum on August 6, he reviewed the IPG proposal and the project's relationship to the NSF. "What I'm really hoping," Feiereisen sums up, "is that these new thoughts about how you do computing could be the catalyst for making sure that Ames and NAS continue to be one of the most exciting places in the world to work."

NASnews
▶ Main Menu
September - October 1998 • Volume 3, Number 5

Information Power Grid Proposal Wins Praise

by Wade Roush

Imagine a system that can transform coal, gas, sunlight, wind, rushing water, and geothermal energy into moving electrons, and then channel those electrons to any location on the continent. This system already exists, of course-it's called the electrical power grid. Now imagine a similar system that can collect, store, and analyze vast amounts of scientific data about the Earth, the air, and outer space, as well as the vehicles that roam them, then transmit, transform, or portray that data for any aerospace engineer who needs it-seamlessly and on demand. This system doesn't exist, but a prototype will before 2000. At NASA, it's called the Information Power Grid (IPG).

The IPG, envisioned as a unified collection of geographically dispersed supercomputers, storage devices, scientific instruments, workstations, and advanced user interfaces, moved a step closer to reality in June. That was when a technical review committee, convening at Ames Research Center, endorsed a preliminary project proposal drafted by a team from NASA's research centers and the National Science Foundation's (NSF) Partnership for Advanced Computing Infrastructure.

Members of the review committee expressed enthusiasm about the idea of "linking NASA's vast resources together to create an intelligent, scalable, adaptive computational capability," as stated in the proposal's introduction. "It's just the right moment to start these sorts of activities," says Ian Foster, a senior scientist in the Mathematics and Computer Science Division at Argonne National Laboratory and a member of the committee. "We face these problems [in aerospace engineering and other fields] that are really too complex...for a single person or a small group sitting at a single location to try and address. That's what the grid is about."

National Organization Needed

The committee of outside experts, including computer scientists and engineers from aerospace companies, other government agencies, and university supercomputing centers, was also receptive to the IPG team's argument that NASA, and perhaps the NAS Systems Division in particular, is uniquely qualified to coordinate the grid's construction. "The program has to be organized on a national basis," says Ray Cosner, manager of computational fluid dynamics (CFD) applications at The Boeing Company in St. Louis and a member of the review committee. "I don't see anybody except NASA, or at least some national organization, that can handle this."

Still, the committee's endorsement was merely a first step toward building the IPG. "They said it's a good idea, but it's just a start," says William Feiereisen, newly appointed chief of the NAS Systems Division (*see article, this issue*). The committee asked the Ames team to write a more detailed project plan by the end of September and to prepare for a Program Readiness Review later in the fall.

Lofty Vision, Down-to-earth Tasks

Ten years ago, no one could have accurately predicted the unprecedented growth of the global Internet. Similarly, the IPG project's main proponents at Ames, including Feiereisen, former NAS division chief F. Ron Bailey, and Ames Aeronautical Information Technologies chief William Van Dalsem, cannot foretell the precise scope and functions of the future IPG; the only certainty is that something like it is needed to assure continued U.S. leadership in high-performance computing for aerospace applications.

The team's proposal describes the IPG as "a geographically distributed heterogeneous computational capability" that will provide "ubiquitous and uniform access to a wide range of computational, communication, data analysis, and storage resources-many of which are specialized and cannot be replicated at all user sites...[it] will link scientists, engineers, design teams, manufacturers, and suppliers into 'virtual enterprises' with access to shared data, resources, facilities, and services."

Behind this lofty vision, however, lie three concrete tasks, according to Feiereisen: "To build a prototype grid, to institute a series of research projects which address the [software] questions that we really don't know how to do, and to pick out a set of applications to test things out, to see whether this concept actually works." The IPG team's proposal calls for delivery of an experimental grid, capable of handling numerically intensive applications flexibly across geographically dispersed locations, by October 1999.

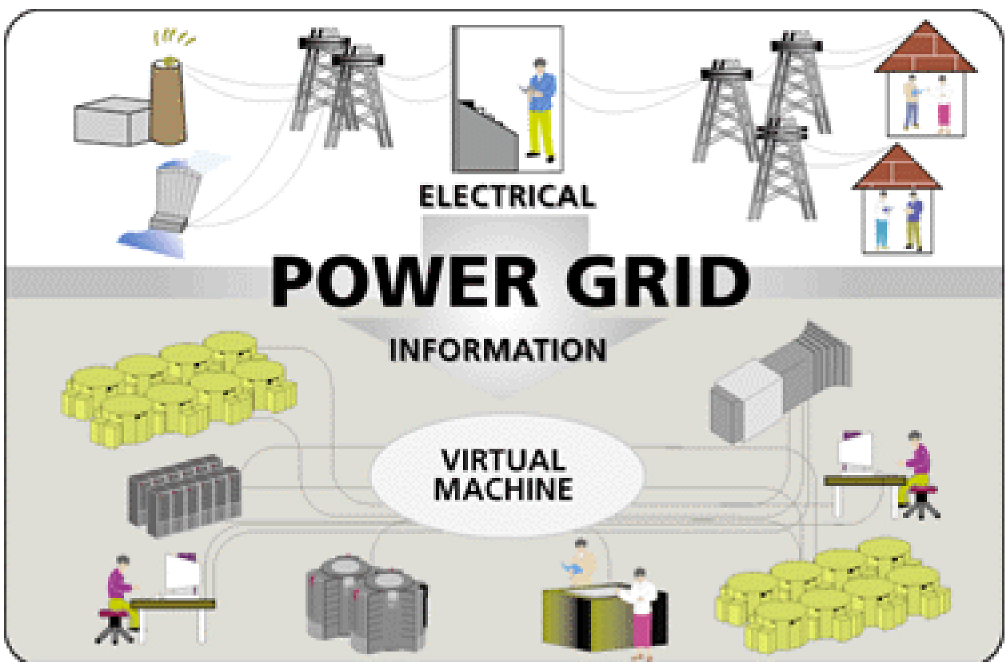

Just as the electrical power grid makes electricity ubiquitous, the Information Power Grid will put vast computational power, storage capacity, and networking capabilities at the fingertips of researchers around the country. The "virtual machine" or "middleware" functions like a control room, matching shifting computational demand with available resources. Graphic by Eunah Choi.

Another task is to build the staff to support the project. Feiereisen has filled two key positions in the IPG management organization. Bill Johnston, IPG project manager, comes from Lawrence Berkeley

Laboratory and has worked on high-speed, data-intensive national networks for the past 10 years. He is the architect of the Department of Energy's grid project called "China Clipper." Dennis Gannon, on leave from Indiana University at Bloomington, where he heads the Computer Science department, is managing the enabling technologies research of the IPG project.

Partners Help IPG Get Up to Speed

Beginning construction of the prototype grid may be the easiest part, since two other organizations with a year's head start in the area of distributed heterogeneous computing have offered to share their expertise and equipment with NASA. In 1997, the NSF PACI program chose two main groups as funding recipients. One is led by the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign. The other is led by the San Diego Supercomputer Center (SDSC) at the University of California, San Diego.

The NCSA-led group, called the National Computational Science Alliance includes more than 50 federal, academic, and industry partners, and is working to build what it calls the National Technology Grid. SDSC spearheads another coalition of several dozen partners known as the National Partnership for Advanced Computing Infrastructure (NPACI) . Feiereisen says that NASA has already signed a formal memorandum of understanding with the NSF to fund research at both the Alliance and NPACI. In return "we get access to their web of intellectual capital instantly," he explains.

'Fusion, Not Competition'

IPG research projects will build on the NSF centers' work in a broad spectrum of areas. One example is middleware-network management software that treats supercomputers and storage devices at many sites as a single "virtual machine" or "metacomputer" (NAS News, *March-April '98*). Alliance members have already built a testbed called GUSTO (co-developed by Argonne's Foster and Carl Kesselman, of the University of Southern California's Information Sciences Institute), which links high-performance computers at 27 sites.

Dave DiNucci, a former member of the NAS parallel tools team who played a significant role in developing the IPG proposal, believes researchers at NAS and other NASA supercomputing centers will be able to learn from the GUSTO experiment and then help scale it up. "We see the PACI programs coming up with ideas, then coming to a place like NAS for integration, testing, and demonstration," he says. "There will be fusion, not competition."

The final task on Feiereisen's implementation plan, identifying test applications for the Information Power Grid, is an area in which NASA is traditionally strong. In addition to aeronautics applications, "the IPG project will have applications related to all the NASA missions," says Feiereisen, "including space sciences, earth sciences, and even life sciences."

Linking Sites, Starting Research

According to the IPG team's proposal, one of the first steps in the projected seven-year, $312 million effort will be to link high-performance computers at NASA's Langley, Lewis, and Ames Research Centers and the NCSA Alliance institutions into a prototype grid. Interfaces will be developed between NAS's PBS (*"Portable Batch Systems Goes Commercial"*) and other candidate components of the "virtual machine," such as the Alliance's Globus system and NPACI's Legion system, among others. To test the prototype grid, researchers will create distributed versions of existing aerospace applications such as Lewis's Numerical Propulsion Simulation System.

The proposal also describes nearly fifty research projects designed to lay the intellectual and technical groundwork for the IPG. Many grow out of research currently under way at Ames and other centers. In the area of network protocols, for example, autonomous reasoning researchers will develop new heuristics, intelligent agents, and neural networks to help system administrators direct the flow of information around the grid with maximum efficiency. Some of this research has already begun, Feiereisen says, and Globus is running on systems at Ames, Langley, and Lewis.

A Revolution in the Making

Ultimately, however, the Ames team's vision for the IPG extends beyond NASA and the NSF PACI centers. "If something like [the IPG] ever does become the model for computing, data access, and collaboration, it could be as exciting as what's happening on the Internet right now," says Feiereisen. "From your home you would be able to do more than just click on somebody's link on a web site. You would be able to query databases in the library of Congress or in the various

NASA space data archives, or you could perform your own simulations. People all across the country would have access to these things."

John Toole, NCSA deputy director for Alliance programs and a member of the IPG proposal review committee, seems to share Feiereisen's excitement. "If NASA takes on the challenges that they have stated," he says "this will be looked on as, I think, a revolutionary time."

NASnews

▶ Main Menu

September - October 1998 • Volume 3, Number 5

In this issue...

DoD Takes Advantage of NAS Metacenter Expertise

by Peter Adams

Within This Article...

The Department of Defense has already consolidated much of its high-performance computing capability into Major Shared Resource Centers (MSRCs) at four sites east of the Mississippi. For several years, however, the department's High Performance Computing Modernization Program has been considering going a step further to implement a "metaqueuing environment" that will allow users to submit and run jobs at whichever MSRC is ideal, regardless of distance. The DoD now plans to get such a project underway, with help from the NAS scientists who built the metacenter linking computing resources at NASA's Ames and Langley Research Centers.

DoD representatives met with NAS personnel at Ames in April to begin drawing up plans for the project. Sharad Gavali, the NAS representative on the joint DoD-NASA executive committee overseeing the project, says the DoD approached NAS because of a compatibility of resources and expertise. DoD supercomputing managers realized that "they did not have to reinvent the wheel," Gavali says. "They could take what [NASA] had done and build upon that."

The goal of a metaqueuing environment is to provide a seamless front-end user interface, joined to an infrastructure which can accept, track, and run users' job submissions. The Langley-Ames Metacenter project began in 1995 and became fully operational in October 1996 (NAS News, *July-August 1997*).

Similar Environments

The NASA metacenter was created on IBM SP2 systems running the Portable Batch System (*see article, this issue*). Of the four DoD MSRCs, just two -- the U.S. Army Corps of Engineers Waterways Experiment Station (CEWES) in Vicksburg, Mississippi, and the Aeronautical Systems Center (ASC) at Wright-Patterson Air Force Base in Ohio -- operated the same environment. These sites were chosen to initiate DoD cross-center computing partly for their similarity to the Ames-Langley architecture, according to Gavali.

Implementation will probably "begin by duplicating what we had done with the SPs at NASA," according to Mary Hultquist, a member of the technical committee working on the project and NAS's group lead for the Ames-Langley Metacenter project. "It's the same setup," she says, with "a few issues specific to the DoD," such as cycle allocations, security authentication, and special projects handling. Incremental technical steps are "progressing well," Hultquist reports, and the project proposal calls for an operational metaqueuing environment by October 1.

Two joint DoD-NASA committees oversee the project. A technical committee consisting of two representatives each from ASC, CEWES, and NAS is studying the feasibility and implementation of different metaqueuing approaches. An executive committee with two representatives each from ASC and CEWES and one from NAS will reconcile the technical committee's recommendations with the specific needs of the DoD sites.

Software Work Continues

While the MSRC integration project gets underway, several NAS groups -- including high speed processing, parallel systems, and scientific consulting -- continue to work on improved job management software that can easily encompass heterogeneous systems. A "heterogeneous peer scheduler," expected to become operational next year, will integrate many different computing resources at each site into the metaqueuing environment. The main challenge in creating such a scheduler, says Hultquist, is the nature of heterogeneous computing itself, since the scheduler must automatically choose the ideal system on which to run a submitted program.

The DoD's current plans encompass only the two MSRCs in Mississippi and Ohio. However, Gavali says, "once the environment is operational, we will work on having other MSRCs join."

The DoD plans to demonstrate the metacomputing concept at November's SC98 conference in Orlando, Florida, using two IBM SP2s at ASC and CEWES.

NASnews

September - October 1998 • Volume 3, Number 5

High-speed Processor Techniques

Life After CF77: You *Can* Get Better Performance

by Art Lazanoff

The Cray Fortran 77 compiler (CF77) was removed from Von Neumann, Eagle, and the Newton cluster last March (*see "Cray Fortran Compiler 'Bytes' the Dust"*), and NAS system users have gradually learned to live with CF90. It was quite unnerving for users to have the familiar, trusted CF77 compiler replaced by the newer, unknown compiler. Yet some users have found that-with some relatively minor changes-they are actually getting better performance with CF90.

'I have to have CF77 back!'

"My code is running twice as slow under CF90" was the complaint from several users. Vectorization and autotasking are integrated into CF90, but the NAS scientific consulting (SciCon) group discovered after investigating several codes that CF90 just doesn't vectorize as well as fpp, the preprocessor in the CF77 compiling subsystem. SciCon group members ran the users' codes through a revived copy of fpp, then fed the output to CF90. The somewhat surprising result: the new compiler then ran the users' codes faster than CF77.

Following this process, the scientific consultants isolated specific code examples and submitted them as Software Problem Reports to Cray Research, or created workarounds using directives to improve performance when CF90 wasn't doing it automatically. Codes that started out running half as fast with CF90 ended up running 10 to 25 percent faster than with CF77.

The NAS scientific consultants use tools such as the Hardware Performance Monitor (HPM) and the performance analyzer perfview to characterize the performance of a code and to compare the CF77 CPU time distribution to the CF90 version. For several users experiencing slow performance, the consultants identified specific code fragments or coding styles, rather than a general performance penalty, as the culprits. SciCon's web pages contain information on how to use these tools, and group members can coach users on how to use them effectively.

User Code Examples
Here's a quick look at several user programs affected by the switch to CF90.

Case 1
In this instance, the user has two large outer loops that surround a number of very short inner loops.

```
SUBROUTINE CASE1(ILEN,JLEN)
PARAMETER (KLEN = 7, LLEN = 3)
DO I = 1, ILEN ! level 1
...
DO J = 1, JLEN ! level 2
...
DO K = 1, KLEN ! level 3
...
DO L = 1, LLEN ! level 4
```

The preprocessor unrolled the loops at levels 3 and 4 (DO K and DO L) and then vectorized the level 2 loop. Since the iteration count was known at compile time, fpp actually wrote the number of iterations out completely, leaving no surviving inner loops at levels 3 and 4. In contrast, CF90 unrolled the level 4 loop and then vectorized the level 3 loop. Since the level 2 loop offered a much longer vector length, CF90 missed the better opportunity. Adding the directives to mimic fpp's approach yields the following:

```
SUBROUTINE CASE1(ILEN,JLEN)
PARAMETER (KLEN = 7, LLEN = 3)
DO I = 1, ILEN ! level 1
...
!DIR$ IVDEP
DO J = 1, JLEN ! level 2
...
!DIR$ NEXTSCALAR
DO K = 1, KLEN ! level 3
...
!DIR$ UNROLL LLEN
DO L = 1, LLEN ! level 4
```

In the same amount of CPU time, this case ran 30 timesteps with CF90 compared to 80 with CF77. With these new directives, CF90 delivers as many as 106 timesteps.

Case 2

In this instance, CF90 vectorized a level 3 loop when it should have unrolled the loop and vectorized the level 2 loop.

```
do 1 m=1,ndp ! level 1
!DIR$ IVDEP
do 1 l=1,ndv(m) ! level 2
np=map1(m,l)
npi=np-1
npj=np-ni
npk=np-ni*nj
!DIR$ UNROLL 5
do 1 il=1,5 ! level 3
y(il,1,np)=x(il,1,npi,1)+x(il,1,npj,2)+x(il,1,npk,3)
y(il,2,np)=x(il,2,npi,1)+x(il,2,npj,2)+x(il,2,npk,3)
y(il,3,np)=x(il,3,npi,1)+x(il,3,npj,2)+x(il,3,npk,3)
y(il,4,np)=x(il,4,npi,1)+x(il,4,npj,2)+x(il,4,npk,3)
y(il,5,np)=x(il,5,npi,1)+x(il,5,npj,2)+x(il,5,npk,3)
1 continue
```

Initially, this code ran 16 percent slower under CF90 than under CF77. The changes produced a code running 12% faster than under CF77.

Case 3
Here, CF90 missed the chance to vectorize the DO loop when it failed to consolidate an IF -THEN-ELSE block.

```
do 1 if=1,linf(linf(ip))
loc=linf(ip)+if
jf=linf(loc)
isurf=lbouf(jf,1)
ibc=bcpch(isurf)
if(ibc .eq. 4) then
nclof=nclof+1
lpoi1(nclof)=jf
lelem(ic)=jf
c
elseif(ibc .eq. 44) then
nclof=nclof+1
lpoi1(nclof)=jf
lelem(ic)=jf
c
elseif(ibc .eq. 55) then
nclof=nclof+1
lpoi1(nclof)=jf
lelem(ic)=jf
c
endif
continue
```

This loop vectorizes as follows:

```
do 1 if=1,linf(linf(ip))
loc=linf(ip)+if
jf=linf(loc)
isurf=lbouf(jf,1)
ibc=bcpch(isurf)
if(ibc .eq. 4 .or. ibc .eq. 44 .or. ibc .eq. 55)
then
nclof=nclof+1
lpoi1(nclof)=jf
lelem(ic)=jf
c
endif
continue
```

Consolidating just one similar IF-THEN-ELSE block was sufficient to get CF90 to run the code in about 190 seconds-compared to 260 with CF77 and 510 on CF90's first try. This user decided that he didn't really need CF77 back after all.

Case 4
In Case 3, above, both CF77 and CF90 missed the opportunity to vectorize code where integer scalar temporaries are used to hold indexes for array references in a loop.

```
do if=1,nclo0
jf=lclof(if)
ip1=lbouf(jf,2)
xp1=coord(ip1,1)
```

These expressions can be rewritten by eliminating the use of the integer scalar temporaries:

```
do if=1,nclo0
xp1=coord(lbouf(lclof(if),2),1)
```

Doing so throughout the code reduced the run-time of Case 3 from 190 to 126 seconds.

SciCon, Cray Work to Help Users
Most calls to the NAS scientific consultants about the new compiler have come from users looking for help converting their CF77 command line options to the CF90 equivalents. A few users have reported bad results or bad performance; in some cases, errors in the users' programs were as the cause of the problem.

Confirmed bugs in CF90 are promptly reported to Cray Research Inc. The SciCon group has already received bug fixes for most of these problems. Cray fixes problems quickly and returns updated, tested compilers on a regular schedule. Dave Dillon, one of the NAS Facility's onsite Cray analysts, provides support to users to identify and isolate bugs and performance problems.

Taking Initial Steps
If you suspect that your code is not running as fast as it should be under CF90, here are some initial steps to follow:

* Use HPM to check your average vector length and to measure your code's MFLOP/s performance.

* Make a perfview run to obtain the CPU time distribution and vector length within each subroutine.

* Obtain a source listing with the options -r2 and -Onegmsgs turned on. This will give you a listing that explains why loops are not vectorized or autotasked. Since perfview will give you a time distribution, you can focus on the top CPU "burners."

Descriptions of HPM and perfview can be found in the NAS Cray Vector Processor User Guide, Chapter 4, "Performance Monitoring Tools for the Von Neumann, Eagle and the Newtons." For more help, contact the NAS scientific consulting group at (650) 604-4444 or (1-800) 331-8737, or send email to nashelp@nas.nasa.gov.

Recent Technical Seminars Available on Video

Many research seminars presented at the NAS Facility are accessible through a videotape loan program. Information on past training events and procedures for obtaining training materials are also available. Visit the NAS Facility's Training web site for more information.

Here are summaries of the most widely attended technical seminars at the NAS Facility since mid-July.

Technologies and Designs for Nanometer-scale Electronic Computers: A Review and Prospectus. At the July 17 Nanotechnology Seminar, James Ellenbogen, principal scientist at The MITRE Corp., discussed the advantages of nanometer-scale electronic computers over other forms of nanocomputers, including the ability to build directly on decades of experience and the high infrastructure of the worldwide electronics industry.

Ellenbogen discussed some of the new problems posed by the fact that nanometer-scale electronic devices (or nanoelectronics) will be only slightly bigger than the individual molecules of which they are comprised. He discussed some of the unique applications of nanoelectronics, along with possible future opportunities for the insertion of nanoelectronics into commercial computer electronics and other high-technology end items.

From Schroedinger to Dirac to Atoms in Molecules. Richard Bader, Department of Chemistry at McMaster University in Canada, has shown that because of an atom's fundamental form, all hypotheses about molecular structure-that is, all of the concepts of chemistry-follow from physics. At the July 28 Nanotechnology Seminar, Bader traced the path of research that lead to this conclusion, with emphasis on how it evolved from a search for the conceptual basis of chemistry in terms of the observable properties of a system.

Contact Determination and Distance Computation: Algorithms, Systems, and Applications. The problems of contact determination and distance computation are ubiquitous within many tasks in computer graphics, virtual environments, robotics, manufacturing, and mechanical design. At the August 6 New Technology Seminar, Ming Lin, Department of Computer Science at the University of North Carolina, gave an overview of work in these areas. Lin also described algorithms and systems that use hierarchical representations composed of higher-order bounding volumes.

AFramework for the Real-time Walkthrough of Massive Models. Also on August 6, Dinesh Manocha, Department of Computer Science at the University of North Carolina, showed how interactive walkthroughs can greatly benefit the multidisciplinary design and review process for CAD applications and scientific visualizations. He gave an overview of recent work on interactive walkthroughs of massive models and presented a framework for rendering very large models at nearly interactive rates.

Experiment Management Support for Parallel Performance Tuning. At the August 11 New Technology Seminar, Karen Karavanic, Computer Sciences Department at the University of Wisconsin, discussed her research on parallel performance tuning. The development of a high-performance parallel system or application is an evolutionary process, with both the code and the environment going through many changes during a program's lifetime.

In Karavanic's research, performance tuning is reframed as a specialized instance of scientific experimentation, and she develops methods for storing, viewing, and using performance data that span a variety of executions, program versions, and environments. She described some early results, as well as work in progress.

High-performance Graphics for the Linux PC. Robert Geist, Department of Computer Science at Clemson University, talked about high-performance graphics for the Linux PC at the August 13 New Technology Seminar. Securing high-performance OpenGL capability has been expensive, and researchers are usually forced to select from a range of workstations offered by Silicon Graphics Inc. Geist described a low-cost, high-performance alternative to such platforms.

A Distributed Version of the SequenceL Language. On August 18, as part of the Information Power Grid series of New Technology Seminars, Daniel Cooke, Computer Science Department at the University of Texas, talked about a distributed version of the SequenceL Language. Current software for the processing and analysis of data is not powerful enough to handle the vast amounts of data that can be acquired and stored by modern hardware-a problem that, according to Cooke, is being discussed at the highest levels of the U.S. government.

Cooke discussed the SequenceL effort, based at the NASA University Research Center at the University of Texas, and how it has lead to commercial interest in the language. He gave an overview of the SequenceL abstraction, focusing on how the abstraction is well suited to the description of "divide-and-conquer-based" distributed data mining solutions.

Robust Simplification Methods for Triangle and Tetrahedral Meshes. One of the most critical and fundamental research problems encountered in the analysis and visualization of massive datasets is the development of efficient methods for storing, approximating, and rendering. According to Kenneth Joy, Center for Image Processing and Integrated Computing aat the University of California, Davis, the problem is to develop different representations of the dataset, where each can be substituted for the complete set, depending on the requirements of the analysis or the visualization technique. Joy presented some new methods for the construction of multiple levels of triangle and tetrahedral meshes.

Advanced Optimization Techniques for High Performance Fortran. On August 21, John Mellor-Crummey, Computer Science Department at Rice University, discussed the Rice dHPF compiler. The project's goal is to develop advanced optimization techniques to provide high performance for a broad spectrum of scientific applications-with minimal restructuring of existing Fortran 77 or Fortran 90 applications. He also described research at Rice that uses dHPF to parallelize the NAS 2.3 serial versions of the benchmarks SP, BT, and LU.

Developing a Scalable Infrastructure for HPSS: First Experience at RUS. At the September 11 New Technology Seminar, Peter Haas, from the Regional Computer Center, University of Stuttgart, Germany, talked about their investigation of new technologies for HSM systems and associated networks. He discussed the pros and cons of DMF, SAMFS, and HPSS and the various techniques needed to "interwork client and server systems at very high speeds."

In this issue...

Credits

Executive Editor: Bill Feiereisen

Editor: Jill Dunbar

Senior Writer: Wade Roush

Contributing Writers: Peter Adams, Art Lazanoff

Image Coordinator: Chris Gong

Special thanks to: Richard Anderson, Ray Cosner, Judy Conlon, James Donald, Ian Foster, Sharad Gavali, Bob Henderson, John Hardman, Mary Hultquist, James P. Jones, Randy Kaemmerer, Leslie Keely, Robin Orans, Marcia Redmond, Leigh Ann Tanner, John Toole, Rita Williams, Sam Uselton

Editorial Board:
Cristy Brickell, Jill Dunbar, Chris Gong, Bill Feiereisen, Nateri Madavan, Patrick Moran, George Myers, Wade Roush, Harry Waddell